

Yet Another Look at the Omitted Variable Bias

Masayuki Hirukawa*
Ryukoku University

Irina Murtazashvili†
Drexel University

Artem Prokhorov‡
University of Sydney and
St.Petersburg State University

September 20, 2020

Abstract

When conducting regression analysis, econometricians often face the situation where some relevant regressors are unavailable in the data set at hand. This paper shows how to construct a new class of nonparametric proxies by combining the original data set with one containing the missing regressors. Imputation of the missing values is done using a nonstandard kernel adapted to mixed data. We derive the asymptotic distribution of the resulting semiparametric two-sample estimator of the parameters of interest and show, using Monte Carlo simulations, that it dominates the solutions involving instrumental variables and other parametric alternatives. An application to the PSID and NLS data illustrates the importance of our estimation approach for empirical research.

Keywords: Data combination; endogeneity; instrumental variable estimation; omitted variable bias; two-sample estimation.

JEL Classification Codes: C13; C14; C31.

*e-mail: hirukawa@econ.ryukoku.ac.jp.

†e-mail: im99@drexel.edu.

‡e-mail: artem.prokhorov@sydney.edu.au.

1 Introduction

Omission of relevant variables leads to challenging problems in applied work. If the correlation between the omitted variable and included regressors is strong, least squares estimates are biased and inconsistent. This issue has a long history in economics.

A classic example of the missing regressor is an ability measure in Mincer’s (1974) wage regression, where estimates suffer from the so called “ability bias” unless the regressors include a variable representing ability (see, e.g., Card, 1995, for details). Micro-level data sets such as the Current Population Survey (CPS) and the Panel Study of Income Dynamics (PSID) do not generally contain individual test scores that can be used as (a proxy for) ability.

Another example can be found in the study of gender wage gap (e.g., Zabalza and Arufat, 1985; Black, Trainor, and Spencer, 1999). Work experience is an important regressor in the wage regression. Although the General Household Survey (GHS) and CPS contain wages and other predictors, they include no actual work experience.

This paper argues that availability of a suitable auxiliary sample offers an under-exploited opportunity. We use the sample to construct a new semiparametric estimator and we study its properties. In a nutshell, our two-sample estimation procedure is to run the ordinary least squares (OLS) after replacing the missing regressor with a nonparametric estimate of its conditional mean, obtained using a new mixed-data kernel on the other sample.

Our approach is different from but not unrelated to the work on matched-sample estimators (see, e.g., Abadie and Imbens, 2006, 2011; Hirukawa and Prokhorov, 2018). One key advantage of our estimator is that the imputed proxy does not rely on a combined sample constructed via the nearest-neighbor matching (NNM) or another ad hoc matching scheme. Matching-free estimation is also the purpose of some well-known moment-based two-sample methods (e.g., Klevmarken, 1982; Angrist and Krueger, 1992, 1995; Arellano and Meghir, 1992; Inoue and Solon, 2010; Murtazashvili et al., 2015). Unfortunately, these approaches are not applicable in the setting of a linear regression where some regressors are missing

in the primary sample, and two-sample moment-based estimation is infeasible. Our setup is closer to what is considered by Pacini (2019). However, his main focus is on relaxing identification conditions for two-sample estimation where missing regressors are imputed from the auxiliary sample.

We establish large-sample properties of the estimator we propose. Its finite-sample properties are examined through extensive Monte Carlo simulations, where we show that our estimator dominates the competitors, in particular, the linear and nonlinear instrumental variable (IV) estimation and the parametric two-sample alternative. In an empirical application, we look for a setting where we are able to obtain estimates based on both the IV estimator and proposed estimator at the same time. We do so for the return to schooling and offer new insights arising from the use of IVs based on one sample as well as from the use of proxies based on two samples.

The nonparametric imputation we propose is important in several practical dimensions. First, having access to a second sample, empirical economists quite commonly impute the variables that are necessary for estimation but missing in the first sample. For example, Fang, Keane, and Silverman (2008) impute medical expenditure of Health and Retirement Study (HRS) respondents using the information from the Medicare Current Beneficiary Survey (MCBS), and Flavin and Nakagawa (2008) impute the square footage of the homes of PSID respondents using data on housing from the American Housing Survey (AHS). Instead of estimating the conditional mean of the missing variable nonparametrically, these papers adopt a linear projection-based imputation. As a matter of course, if the functional form turns out to be incorrectly specified, then consistency of the estimators can no longer be guaranteed. Furthermore, estimation errors associated with imputation need to be explicitly accounted for, as imputed values cannot be treated as error-free and adjustments need to be made to the standard errors. A remarkable observation is that our nonparametric imputation approach performs no worse even under correct parametric specification. The result on inconsistency of the linear projection approach and the circumstances under which it still works may be of independent interest.

Second, our nonparametric imputation method is important in the context of the literature on two-step semiparametric regression estimation with a plug-in first-step nonpara-

metrically generated regressor (see, e.g., Rilstone, 1996; Stengos and Yan, 2001; Newey, 2009; Hsu, Liao and Lin, 2018). The literature explores statistical properties of regression estimators when the regressors that cannot be observed directly are estimated by nonparametric methods, similarly to Pagan’s (1984) classical fully parametric case. However, these approaches are not directly applicable when imputation needs to be based on two samples or when there are no common observations in them.

Finally, we employ a new uniform consistency result for estimators based on a nonstandard kernel and mixed data, which is of independent interest. A key aspect of the result is compactness of the support of the kernel function, a feature that often arises in economics and finance either by construction or as a theoretical requirement. For example, such economic and financial variables as expenditure and budget shares, unemployment rates and default and recovery rates are typically bounded from above and below by construction. Theoretical compactness requirements are used in the context of partial linear regression models (see, e.g., Yatchew, 1997), the NNM (see, e.g., Abadie and Imbens, 2006), and estimation of first-price auctions (see, e.g., Guerre, Perrigne and Vuong, 2000), to name just a few examples. Interested readers may consult Hirukawa, Murtazashvili and Prokhorov (2020) for further details.

The remainder of this paper is organized as follows. Section 2 describes the model, defines the proposed estimator, derives and discusses its convergence properties. Section 3 presents results of selected Monte Carlo simulations to compare finite-sample properties of the proposed estimator with the IV and other alternatives. As an empirical example, in Section 4, we apply the two-sample estimator to a version of Mincer’s (1974) wage regression. Both in simulations and application, we compare our estimator (a) with the IV approach, (b) with an approach where the missing variable is imputed using a fully parametric method and (c) with the approach of Hirukawa and Prokhorov (2018) based on NNM. Section 5 concludes with a few questions for future research. Technical proofs are provided in the Appendix. Moreover, the online Supplement contains details of the alternative estimators (a)-(c), comprehensive simulation results for a variety of designs, and additional discussions on the empirical example.¹

¹Available at <https://sites.google.com/site/artembprokhorovv0/papers/Supplement.pdf>

The paper adopts the following notational conventions: $\|A\| = \{\text{tr}(A'A)\}^{1/2}$ is the Euclidean norm of matrix A ; $\mathbf{1}\{\cdot\}$ denotes an indicator function; $0_{p \times q}$ signifies the $p \times q$ zero matrix, where the subscript may be suppressed if $q = 1$; $B(p, q) = \int_0^1 y^{p-1} (1-y)^{q-1} dy$ for $p, q > 0$ is the beta function; and the symbol $>$ applied to matrices means positive definiteness.

2 Two-Sample Two-Step Estimation

2.1 Model

We consider the following linear regression model

$$Y = \beta_0 + X_1' \beta_1 + X_2' \beta_2 + X_{3I}' \beta_3 + u, \quad (1)$$

where $X_1 \in \mathbb{R}^{d_1}$, $X_2 \in \mathbb{R}^{d_2}$, $X_{3I} \in \mathbb{R}^{d_{3I}}$ (the distinction between these regressors will be made clear shortly). Throughout we assume that either β_1 or β_3 is the parameter of main interest. Let $d := d_1 + d_2 + d_{3I}$. When $(1, X_1', X_2', X_{3I}')' \in \mathbb{R}^{d+1}$ is exogenous and a single random sample of (Y, X_1, X_2, X_{3I}) is available, the OLS estimator of $\beta = (\beta_0, \beta_1', \beta_2', \beta_3')'$ is consistent under the usual assumptions.

Let \mathcal{S}_1 denote the data set at hand and let \mathcal{S}_2 denote a second data set which will be required for a two-sample estimation. We assume that $\mathcal{S}_1 = (Y, X_1, X_3)$ and $\mathcal{S}_2 = (X_2, X_3)$, where $X_3 := (X_{3I}, X_{3E})$ is the vector of common variables across the two samples that are partitioned into those included (X_{3I}) and those excluded (X_{3E}) from regression (1). We note that X_2 is missing in \mathcal{S}_1 although it is assumed to be relevant, i.e., $\beta_2 \neq 0$. This setup was recently considered by Pacini (2019) but his focus was on relaxing the identification assumptions involving the joint distribution of (Y, X_1, X_2, X_{3I}) .

The distinction between X_1 and X_3 will become important later. For now, we can think of X_3 as variables that are common across \mathcal{S}_1 and \mathcal{S}_2 , enter as regressors and are used, due to their presence in both samples, for imputation of the missing regressors in X_2 . On the other hand, when a common variable serves merely as a regressor but is not used for imputation, it is classified as X_1 . The vector X_1 can be empty; the vector X_3 cannot.

Even though \mathcal{S}_1 and \mathcal{S}_2 contain common variables X_3 , this does not mean they need to have common observations, i.e., the two samples do not have to overlap. Finally, denote

$d_3 := \dim(X_3) = \dim(X_{3I}) + \dim(X_{3E}) := d_{3I} + d_{3E}$, where $d_3 > 0$ must be the case, and either d_{3I} or d_{3E} is allowed to be zero.

An example of this sampling arrangement is when a micro-level data set, e.g., CPS or PSID, can be complemented by an auxiliary data set which is focused on some important variables that are not available elsewhere, e.g., a psychometric data set with ability measures.

In such cases it is common to instrument for the unobserved variable (see Supplement for details of IV estimation). However, the IV estimation is unavailable if there are no valid instruments in the same sample and it may be desirable to avoid the IV estimation for a host of other reasons discussed in Section 2.5.

Denote $\mathcal{S}_1 = \mathcal{S}_{1n} = \{(Y_i, X_{1i}, X_{3i})\}_{i=1}^n$ and $\mathcal{S}_2 = \mathcal{S}_{2m} = \{(X_{2j}, X_{3j})\}_{j=1}^m$. We now discuss how we construct a proxy for X_2 using \mathcal{S}_{2m} .

2.2 Additional Notation

We start by introducing some additional notation that will help define the two-sample semiparametric estimator and derive its asymptotic properties. First, we allow for the vector of common variables $X_3 = (X_{3I}, X_{3E})$ to consist of both continuous (C) and discrete (D) variables so that $d_3 = \dim(X_3) = \dim(X_{3C}) + \dim(X_{3D}) =: d_{3C} + d_{3D}$. While $d_{3D} = 0$ is allowed, $d_{3C} > 0$ must be the case for the subsequent asymptotic analysis. This distinction will be needed for the application of kernel smoothing to mixed continuous and discrete data. Second, let $\mathbb{X}_3 := \mathbb{X}_{3C} \times \mathbb{X}_{3D}$, where $\mathbb{X}_{3C} := \text{supp}(X_{3C})$ and $\mathbb{X}_{3D} := \text{supp}(X_{3D})$. Third, denote $g_2(X_3) := E(X_2|X_3)$ and $\eta_2 := X_2 - g_2(X_3)$ with $E(\eta_2|X_3) = 0$. The conditional expectation function $g_2(\cdot)$ and the error term η_2 will be determining the asymptotic adjustments required to account for the use of nonparametric proxies. Finally, for a pair of vectors (Z, W) , let $\Phi_{Z,W} := E(Z_i W_i')$, $S_{n,Z,W} := (1/n) \sum_{i=1}^n Z_i W_i'$ and $S_{m,Z,W} := (1/m) \sum_{j=1}^m Z_j W_j'$, where summations for $S_{n,Z,W}$ and $S_{m,Z,W}$ are assumed to be taken within \mathcal{S}_{1n} and \mathcal{S}_{2m} , respectively.

2.3 Estimation of β and $g_2(\cdot)$

We reformulate the long regression (1) as

$$Y = X'\beta + \epsilon,$$

where $X := (1, X'_1, g_2(X_3)', X'_{3I})'$ and $\epsilon := u + \eta'_2\beta_2$.

In essence, we wish to obtain an estimator of the unknown function $g_2(X_3)$ using \mathcal{S}_2 and construct a proxy for X_2 using some distance measure between the common variables X_3 in \mathcal{S}_1 and \mathcal{S}_2 . We will use a kernel function (as opposed to NNM or another matching scheme) to do that.

Let $\hat{g}_2(\cdot)$ be some consistent nonparametric estimator of $g_2(\cdot)$. Then, the estimation procedure we propose has the following two steps:

Step 1: Regard $\{(X_{2j}, X_{3j})\}_{j=1}^m$ in \mathcal{S}_2 and $\{X_{3i}\}_{i=1}^n$ in \mathcal{S}_1 as m data points and n design points, respectively, and obtain n nonparametric estimates $\{\hat{g}_2(X_{3i})\}_{i=1}^n$.

Step 2: Run OLS for the regression of Y_i on $\hat{X}_i := (1, X'_{1i}, \hat{g}_2(X_{3i})', X'_{3Ii})'$.

The estimator of β is then

$$\hat{\beta}_{PILS} := S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, Y} = \left(\frac{1}{n} \sum_{i=1}^n \hat{X}_i \hat{X}'_i \right)^{-1} \frac{1}{n} \sum_{i=1}^n \hat{X}_i Y_i.$$

Because this estimator is the OLS with $\hat{g}_2(X_{3i})$ plugged in place of the missing regressor X_{2i} , we call it the *plug-in least squares* (PILS) estimator hereinafter.

At a first glance, this approach – just like its fully parametric counterpart discussed in the online Supplement – may be viewed as a variant of generated regressors, which have been extensively studied by many researchers. However, two samples and nonparametric imputation raise new and nontrivial issues, which require careful consideration. We study them in the next section.

A remaining task is to deliver a consistent estimator of $g_2(\cdot)$. Taking into consideration that $g_2(\cdot)$ may depend on both continuous and discrete covariates, we choose the kernel regression smoother for mixed continuous and categorical data proposed by Racine and Li (2004). In this sense, PILS may be viewed as a partial mean estimator of Newey (1994),

although Newey (1994) does not consider the situation where design and data points for kernel smoothing come from different data sources.

A key part of the estimator is the construction of a multivariate kernel. Let $\mathcal{K}(t_{3C}; x_{3C}, \mathbf{h})$ and $\mathcal{L}(t_{3D}; x_{3D}, \lambda)$ denote product kernels for the continuous and discrete components of X_3 , respectively, where t_3 and x_3 are data points and design points, and \mathbf{h} and λ are vectors of bandwidths. We provide details of how these kernels are constructed in Appendix A.1. Then, the product kernel for X_3 is

$$\mathcal{W}(t_3; x_3, \mathbf{h}, \lambda) := \mathcal{K}(t_{3C}; x_{3C}, \mathbf{h}) \mathcal{L}(t_{3D}; x_{3D}, \lambda).$$

It follows that a nonparametric estimator of $g_2(\cdot)$ can be defined as

$$\hat{g}_2(X_{3i}) := \frac{\sum_{j=1}^m X_{2j} \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)}{\sum_{j=1}^m \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)}, \quad i = 1, \dots, n.$$

There are many conventional options for the specific univariate kernel to use for continuous variables due to the requirement for compactness of \mathbb{X}_{3C} (see Assumption 2 below). Assuming without loss of generality that the compact set is a d_{3C} -dimensional unit hypercube $[0, 1]^{d_{3C}}$, we propose to employ Chen's (1999) univariate beta kernel with support on $[0, 1]$ as an attractive alternative to standard symmetric kernels. Our application of this kernel is based on a theoretical convergence result on a compact support which we prove in a companion paper (Hirukawa, Murtazashvili and Prokhorov, 2020) and which is, as far as we know, new. This asymmetric kernel is free of the boundary bias by construction and its shape varies across design points even under a fixed value of the smoothing parameter b . The latter property implies that the amount of smoothing by this kernel changes in an adaptive manner.

2.4 Convergence Properties of PILS

2.4.1 Regularity Conditions

Now we explore asymptotic properties of the PILS estimator as both n and m diverge. For this purpose, the following regularity conditions are imposed.

Assumption 1. The two random samples $(\mathcal{S}_1, \mathcal{S}_2) = (\mathcal{S}_{1n}, \mathcal{S}_{2m})$ are drawn independently from the joint distribution of (Y, X_1, X_2, X_3) with finite fourth-order moments.

Assumption 2. X_{3C} is continuously distributed with a convex and compact support \mathbb{X}_{3C} , and its density is bounded and bounded away from zero on \mathbb{X}_{3C} .

Assumption 3.

(i) $E(u|X_1, X_3) = 0$ and $\sigma_u^2(X_1, X_3) := E(u^2|X_1, X_3) \in (0, \infty)$.

(ii) $(\eta_2 \perp\!\!\!\perp X_1)|X_3$.

(iii) $g_2(\cdot)$ is non-constant on \mathbb{X}_{3C} if X_{3E} contains at least one continuous variable, and $g_2(\cdot)$ is strictly nonlinear on \mathbb{X}_{3C} otherwise.

Assumption 4.

(i) Let $f(\cdot)$ be the marginal pdf of X_{3C} . Then, for some integer $\nu \geq 2$, the ν th-order derivatives of $f(\cdot)$ and $f(\cdot)g_2(\cdot)$ with respect to X_{3C} are continuous and bounded uniformly on \mathbb{X}_{3C} .

(ii) $\Phi_{X,X} := E(XX') > 0$.

(iii) There exist some constants $\gamma \in (0, \infty)$ and $C \in [1, \infty)$ so that

$$\sup_{x_3 \in \mathbb{X}_3} E\left(\|X_2\|^{2+\gamma} \mid X_3 = x_3\right) \leq C.$$

Assumption 5. The univariate continuous kernel K is either: (a) a symmetric and bounded kernel function of order ν that satisfies (i) $\int K(t) dt = 1$, $\int t^l K(t) dt = 0$ for $l = 1, \dots, \nu-1$, $\int t^\nu K(t) dt \neq 0$, (ii) the first-order Lipschitz condition, and (iii) $\int |t|^k |K(t)| dt < \infty$ for some $k > 2\nu$; or (b) the beta kernel of Chen (1999).

Assumption 6. Sequences of smoothing parameters $h_p (= h_p(m) > 0)$, $b_p (= b_p(m) > 0)$ and $\lambda_q (= \lambda_q(m) \in (0, 1))$ and boundary parameters $\theta_p (= \theta_p(m) > 0)$ for $p = 1, \dots, d_{3C}$ and $q = 1, \dots, d_{3D}$ satisfy one of the following conditions as $m \rightarrow \infty$: (a) for a symmetric ν th-order kernel, $h_p, \lambda_q \rightarrow 0$ and $\log m / \left(m \prod_{p=1}^{d_{3C}} h_p\right) \rightarrow 0$; and (b) for the beta kernel, $b_p, \lambda_q, \theta_p \rightarrow 0$, $b_p/\theta_p \rightarrow 0$ and $\log m / \left(m \sqrt{\prod_{p=1}^{d_{3C}} b_p \theta_p}\right) \rightarrow 0$.

The first three assumptions are often encountered in matching-based estimation (see, e.g., Hirukawa and Prokhorov, 2018) and are natural in our settings. In particular, it follows from Assumption 3(i) that even when some part of X_3 does not enter regression (1) due to an exclusion restriction, X_1 and X_3 are exogenous in the full model. Furthermore, Assumption 3(ii) implies that $\eta_2|X_1, X_3 \stackrel{d}{=} \eta_2|X_3$, where “ $\stackrel{d}{=}$ ” denotes equality in distribution. It seems difficult to avoid the conditional independence or a similar condition (e.g., Assumption c of Inoue and Solon, 2010; Assumption 4 of Hirukawa and Prokhorov, 2018); indeed, this condition helps derive the asymptotic variance of our estimator in a straightforward manner and makes the associated covariance estimation fully operational. Finally, nonlinearity of $g_2(\cdot)$ in Assumption 3(iii) is required only when all continuous common variables are included as regressors in (1). Otherwise, excluded continuous common variables introduce additional randomness in $g_2(\cdot)$, which helps identify β .

Along with compactness of \mathbb{X}_{3C} in Assumption 2, Assumptions 4–6 are required to establish weak uniform consistency of the nonparametric regression estimator $\hat{g}_2(\cdot)$ on \mathbb{X}_3 . Similar conditions can be found, for instance, in Li and Ouyang (2005), Hansen (2008) and Su, Murtazashvili and Ullah (2013) for a symmetric kernel, and Hirukawa, Murtazashvili and Prokhorov (2020) for the beta kernel. Notice that the beta kernel is nonnegative so that the order $\nu = 2$ is the case in Assumption 4(i).

Strictly speaking, uniform consistency holds on \mathbb{X}_3 for a symmetric kernel. In contrast, for the beta kernel, the boundary parameters $\theta_1, \dots, \theta_{d_{3C}}$ play a key role in uniform consistency of $\hat{g}_2(\cdot)$. Specifically, uniform consistency in this case is established on $\mathbb{S}_{X_{3C}} \times \mathbb{X}_{3D}$, where $\mathbb{S}_{X_{3C}} := \prod_{p=1}^{d_{3C}} [\theta_p, 1 - \theta_p]$ is a compact set expanding to $\mathbb{X}_{3C} = [0, 1]^{d_{3C}}$ (more slowly than the smoothing parameters $b_1, \dots, b_{d_{3C}}$) so that $\mathbb{S}_{X_{3C}} \times \mathbb{X}_{3D}$ is expanding to \mathbb{X}_3 as $m \rightarrow \infty$. On the other hand, because the regression estimator $\hat{g}_2(\cdot)$ is a Nadaraya-Watson-type estimator, it suffers from the so-called boundary bias when a symmetric kernel is employed. When the beta kernel is used, $\hat{g}_2(\cdot)$ is free of this issue by construction.

Also observe that Assumption 5(a) allows for higher-order kernels (i.e., $\nu > 2$). In theory, such kernels could be a remedy for the curse of dimensionality implied by Corollary 1 below. In practice, however, it is well known that higher-order kernels are unlikely to dominate nonnegative ones for moderate sample sizes, and that the advantage, if any, is typically

marginal (see, e.g., Marron and Wand, 1992; Marron, 1994). In fact, our simulation results indicate poor finite-sample performance of PILS using higher-order kernels.

2.4.2 Consistency and Asymptotic Normality of PILS

The next two theorems establish consistency and asymptotic normality of $\hat{\beta}_{PILS}$. In particular, asymptotic normality of $\hat{\beta}_{PILS}$ is obtained after correcting for B_{g_2} , the bias term due to kernel smoothing. Each theorem holds regardless of the number of common variables and regardless of the divergence patterns in (n, m) .

Theorem 1. *If Assumptions 1–6 hold, then $\hat{\beta}_{PILS} \xrightarrow{p} \beta$ as $n, m \rightarrow \infty$.*

Theorem 2. *If Assumptions 1–6 hold, then*

$$\sqrt{n} \left(\hat{\beta}_{PILS} - \beta - B_{g_2} \right) \xrightarrow{d} N \left(0_{(d+1) \times 1}, V \right) := N \left(0_{(d+1) \times 1}, \Phi_{X,X}^{-1} \Omega \Phi_{X,X}^{-1} \right)$$

as $n, m \rightarrow \infty$, where

$$\begin{aligned} B_{g_2} &:= S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, [g_2(X_3) - E\{\hat{g}_2(X_3)|X_3\}]' \beta_2} \\ &= \left(\frac{1}{n} \sum_{i=1}^n \hat{X}_i \hat{X}_i' \right)^{-1} \frac{1}{n} \sum_{i=1}^n \hat{X}_i [g_2(X_{3i}) - E\{\hat{g}_2(X_{3i})|X_{3i}\}]' \beta_2 \\ &= \begin{cases} O_p \left(\sum_{p=1}^{d_{3C}} h_p^\nu + \sum_{q=1}^{d_{3D}} \lambda_q \right) & \text{for a symmetric } \nu\text{-th-order kernel} \\ O_p \left(\sum_{p=1}^{d_{3C}} b_p + \sum_{q=1}^{d_{3D}} \lambda_q \right) & \text{for the beta kernel} \end{cases}, \text{ and} \\ \Omega &:= \Omega_1 + \Omega_2 := E \left(X X' \epsilon^2 \right) + E \left(X X' \beta_2' \eta_2 \eta_2' \beta_2 \right). \end{aligned}$$

As derived in the proof of Theorem 1, there is the asymptotically negligible bias term B_{g_2} , due to kernel smoothing. There are also two \sqrt{n} -asymptotically normal terms. One comes from the sampling error in regression (1), and it has the asymptotic variance $\Phi_{X,X}^{-1} \Omega_1 \Phi_{X,X}^{-1}$. The other is due to the approximation error of g_2 , and its asymptotic variance is $\Phi_{X,X}^{-1} \Omega_2 \Phi_{X,X}^{-1}$. As a consequence, the asymptotic variance Ω in Theorem 2 has similar structure to that of the double kernel estimator by Stengos and Yan (2001, Theorem 1) and the two-step series estimator by Hsu, Liao and Lin (2018, Theorem 2.1). A difference from Stengos and Yan (2001) and Hsu, Liao and Lin (2018) is that the two terms are uncorrelated. Each of them can be obtained within two independent samples \mathcal{S}_1 and \mathcal{S}_2 , and thus Ω is free of covariance terms.

2.4.3 Covariance Estimation

Covariance estimation is essential for inference. The problem of estimating V consistently boils down to proposing consistent estimators of Ω_1 and Ω_2 . For the PILS residual $\hat{\epsilon}_i := Y_i - \hat{X}_i' \hat{\beta}_{PILS}$, an estimator of Ω_1 is

$$\hat{\Omega}_1 := \frac{1}{n} \sum_{i=1}^n \hat{X}_i \hat{X}_i' \hat{\epsilon}_i^2.$$

Consistency of $\hat{\Omega}_1$ can be established in the same manner as in the proofs of Theorems 1 and 2.

On the other hand, $\hat{\Omega}_2 := (1/n) \sum_{i=1}^n \hat{X}_i \hat{X}_i' \hat{\beta}_{2,PILS}' \hat{\eta}_{2i} \hat{\eta}_{2i}' \hat{\beta}_{2,PILS}$ might work as a consistent estimator for Ω_2 , where $\hat{\beta}_{2,PILS}$ is the PILS estimator of β_2 and $\hat{\eta}_2 := X_2 - \hat{g}_2(X_3)$. In reality, it turns out to be difficult to obtain $\{\hat{\eta}_{2i}\}_{i=1}^n$ because of the absence of X_2 in \mathcal{S}_1 . Instead, using $\eta_2 | X_1, X_3 \stackrel{d}{=} \eta_2 | X_3$ yields

$$\Omega_2 = E \left\{ X X' \beta_2' E \left(\eta_2 \eta_2' \middle| X_3 \right) \beta_2 \right\} =: E \left\{ X X' \beta_2' \Sigma_2(X_3) \beta_2 \right\}.$$

Then, we can alternatively consider the estimator

$$\hat{\Omega}_2 := \frac{1}{n} \sum_{i=1}^n \hat{X}_i \hat{X}_i' \hat{\beta}_{2,PILS}' \hat{\Sigma}_2(X_{3i}) \hat{\beta}_{2,PILS},$$

where

$$\hat{\Sigma}_2(X_{3i}) := \frac{\sum_{j=1}^m \hat{\eta}_{2j} \hat{\eta}_{2j}' \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)}{\sum_{j=1}^m \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)}, \quad i = 1, \dots, n,$$

and the nonparametric regression residuals $\{\hat{\eta}_{2j}\}_{j=1}^m = \{X_{2j} - \hat{g}_2(X_{3j})\}_{j=1}^m$ can be obtained within \mathcal{S}_2 . The estimation procedure of $\hat{\Sigma}_2(\cdot)$ is inspired by the Algorithm in Section 2.4 of Fan and Yao (1998). To show consistency of $\hat{\Omega}_2$, we can see that a similar argument to the proof of Lemma A1 of Appendix A.2 establishes uniform consistency of $\hat{\eta}_{2j}$ for η_{2j} , which in turn leads to uniform consistency of $\hat{\Sigma}_2(X_{3i})$ for $\Sigma_2(X_{3i})$. Then, $\hat{\Omega}_2 = (1/n) \sum_{i=1}^n X_i X_i' \beta_2' \Sigma_2(X_{3i}) \beta_2 + o_p(1) \xrightarrow{p} \Omega_2$.

In conclusion, V can be consistently estimated by

$$\hat{V} := S_{n, \hat{X}, \hat{X}}^{-1} \hat{\Omega} S_{n, \hat{X}, \hat{X}}^{-1} := S_{n, \hat{X}, \hat{X}}^{-1} \left(\hat{\Omega}_1 + \hat{\Omega}_2 \right) S_{n, \hat{X}, \hat{X}}^{-1}.$$

2.4.4 \sqrt{n} -Consistency of PILS

Theorems 1 and 2 jointly imply that although PILS is consistent, its convergence is affected by the bias term B_{g_2} generated by kernel smoothing. For a large d_{3C} , the order of magnitude in the bias term dominates and the convergence rate of $\hat{\beta}_{PILS}$ becomes inferior. It appears that this problem is unavoidable in a regression that uses a nonparametric component. A similar phenomenon arises in the context of imputation bias correction (see, e.g., Hirukawa and Prokhorov, 2018).

However, \sqrt{n} -consistency of $\hat{\beta}_{PILS}$ automatically holds for small values of d_{3C} . To illustrate such cases, we modify Assumption 6 in order to control the bias and variance convergence of $\hat{g}_2(\cdot)$ more easily.

Assumption 6'. Sequences of the smoothing and boundary parameters satisfy one of the following conditions as $m \rightarrow \infty$: (a) for a ν th-order symmetric kernel, there is a bandwidth $h (= h(m) > 0)$ so that $h_1, \dots, h_{d_{3C}} \propto h$, $\lambda_1, \dots, \lambda_{d_{3D}} \propto h^\nu$, $h \rightarrow 0$, and $\log m / (mh^{d_{3C}}) \rightarrow 0$; and (b) for the beta kernel, there are a smoothing parameter $b (= b(m) > 0)$ and a boundary parameter $\theta (= \theta(m) > 0)$ so that $b_1, \dots, b_{d_{3C}} \propto b$, $\lambda_1, \dots, \lambda_{d_{3D}} \propto b$, $\theta_1, \dots, \theta_{d_{3C}} \propto \theta$, $b, \theta \rightarrow 0$, $b/\theta \rightarrow 0$, and $\log m / \left\{ m (b\theta)^{d_{3C}/2} \right\} \rightarrow 0$.

The corollary below describes the cases in which $\hat{\beta}_{PILS}$ becomes \sqrt{n} -asymptotically normal.

Corollary 1. Let $h \propto (\log m/m)^\alpha$ and $b \propto (\log m/m)^{2\alpha}$ for some constant $\alpha > 0$. Also suppose that one of the following divergence patterns in (n, m) is true: (i) $n/m \rightarrow \kappa \in (0, \infty)$; (ii) $n/m \rightarrow 0$; or (iii) $n/m \rightarrow \infty$ and $n/m^{4\alpha} \rightarrow 0$. Then, under Assumptions 1-5 and 6', as $n, m \rightarrow \infty$, $\sqrt{n} \left(\hat{\beta}_{PILS} - \beta \right) \xrightarrow{d} N(0_{(d+1) \times 1}, V)$ holds either (a) if a ν th-order symmetric kernel is employed and $d_{3C} \leq 2\nu - 1$ or (b) if the beta kernel is employed and $d_{3C} \leq 3$.

Corollary 1 presents an upper bound of the number of continuous common variables for PILS to be \sqrt{n} -consistent. This is recognized as the curse of dimensionality in continuous common variables, and it can be relaxed, in theory, by using higher-order kernels. In

particular, when employing a nonnegative kernel (which is either a symmetric pdf or the beta kernel), we must limit the number of continuous common variables to 3 or less to make PILS \sqrt{n} -consistent. To see this in more detail, note that for the most realistic case $n/m \rightarrow \kappa$, shrinkage rates of the smoothing parameters are $h \propto (\log m/m)^\alpha$ and $b \propto (\log m/m)^{2\alpha}$ for some $\alpha \in (1/4, 1/d_{3C})$ (see Appendix A.4). These rates are faster than what yields Stone's (1982) optimal global rate of convergence, i.e., $h^* \propto (\log m/m)^{1/(4+d_{3C})}$ and $b^* \propto (\log m/m)^{2/(4+d_{3C})}$ that can be obtained through balancing the squared bias and variance of $\hat{g}_2(\cdot)$. This is because to attain \sqrt{n} -consistency we should keep the convergence rate of the dominant bias term in $\hat{g}_2(\cdot)$ sufficiently large by undersmoothing, or more specifically, by setting the exponent α within the aforementioned range. We give some examples of h and b that fulfill the rate requirement in Sections 3 and 4.

Finally, it may be the case that the missing regressor X_2 can be observed with a measurement error. This issue occurs, for example, when an ability measure is unreliable in the wage regression. Suppose that we can observe $\check{X}_2 = X_2 + v$ at best, where v is the measurement error. As long as $E(v|X_3) = 0_{d_2 \times 1}$, the effect of v is filtered out via kernel smoothing and

$$\check{g}_2(X_{3i}) := \frac{\sum_{j=1}^m \check{X}_{2j} \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)}{\sum_{j=1}^m \mathcal{W}(X_{3j}; X_{3i}, \mathbf{h}, \lambda)} = \hat{g}_2(X_{3i}) + o_p(1), \quad i = 1, \dots, n$$

holds. In the end, consistency of PILS is maintained.

2.5 Comparison of PILS with Competing Estimators

We conclude this section by comparing PILS with competing estimators. What we mean by competing estimators is the linear and nonlinear IV, a fully parametric alternative to PILS (called PARA hereinafter), and the matched-sample indirect inference (MSII) and fully-modified MSII (MSII-FM) estimators. PARA is often a method of choice (see, e.g., Fang, Keane, and Silverman, 2008). MSII(-FM) was proposed by Hirukawa and Prokhorov (2018) as an alternative to the inconsistent OLS estimator using the matched sample (called MSOLS hereinafter). Definitions and statistical properties of these estimators are reviewed in the online Supplement. Finite-sample properties of all the estimators will be assessed against PILS in the next section.

2.5.1 Comparison with IV

The use of IVs has been associated with at least two problems. First, instruments that are not strongly correlated with the endogenous variables – weak instruments – lead to large inconsistencies in the IV estimators even if only a weak correlation exists between the instruments and the structural equation error, i.e., even if the instruments only slightly violate the validity assumption (see, e.g., Bound, Jaeger, and Baker, 1995).

When we are ready to assume that the IVs are valid, the weak instrument problem results in large asymptotic and finite-sample biases and, when many such instruments are available, the distribution of IV estimators is known to deviate substantially from the large-sample approximation (see, e.g., Bekker, 1994; Chao and Swanson, 2005).

Second, finite-sample properties of IV estimators are known to be generally poor, especially when the instruments are weak (see, e.g., Flores-Lagunes, 2007). Various bias reduction techniques have had limited success and even very inventive uses of instruments have, for these reasons, been criticized in terms of their validity, strength and finite-sample performance.

Furthermore, valid and strong instruments may be unavailable in the same sample. A number of ingenious methods have been proposed to combine data from more than one source in such cases (see, e.g., Klevmarken, 1982; Angrist and Krueger, 1992, 1995; Arellano and Meghir, 1992; Inoue and Solon, 2010; Murtazashvili et al., 2015). However, the two-sample IV estimators inherit the same problems as their single-sample counterparts (see, e.g., Choi, Gu and Shen, 2018).

The identification strategy is different for our estimator. The central identification conditions are that $g_2(\cdot)$ either depends on additional exogenous variables X_{3E} or is nonlinear. If there are excluded exogenous and relevant variables, then we could use them as IVs. We return to this possibility in the simulations and application. If there are no additional exogenous variables, then identification relies on nonlinearity in the first step. We can exploit this nonlinearity in the first step with IV as we do in the simulation and application. The weak identification issue, similar to the weak IV, can appear in our setting if $g_2(\cdot)$ depends very weakly on X_{3E} , or if $g_2(\cdot)$ is weakly nonlinear. We consider these cases in the extensive

simulation comparisons.

The validity of instruments can be checked. Note that X_2 admits the reduced form $X_2 = g_2(X_3) + \eta_2 = g_2(X_{3I}, X_{3E}) + \eta_2$. If X_{3I} is non-empty (i.e., some common variables are included as regressors) and some elements of Z are relevant, then those elements are either a part of X_{3I} or correlated with X_{3I} . As a result, Z and X_2 are also correlated, and the IV estimator of the omitted variable regression becomes inconsistent. Then, we are tempted to test the null of consistency of IV estimation indirectly by testing the null hypothesis $H_0 : g_2(X_{3I}, X_{3E}) = g_2(X_{3E})$ *a.e.* against the alternative $H_1 : g_2(X_{3I}, X_{3E}) \neq g_2(X_{3E})$ for a non-empty set on $\text{supp}(X_{3I})$. Several versions of the test of significance in nonparametric regressions have been proposed in the literature. Examples include Fan and Li (1996) and Racine (1997) for continuous regressors and Lavergne (2001) and Racine, Hart and Li (2006) for discrete or categorical regressors, to name a few.

2.5.2 Comparison with PARA

PARA imputes a parametric, linear predictor of X_2 . It seems attractive, as it is not subject to the curse of dimensionality in continuous common variables unlike PILS. However, it requires $\dim(X_{3E}) > 0$ for identification because otherwise, the imputed predictor of X_2 is a linear function of X_{3I} . In contrast, as in Assumption 3(iii), PILS achieves identification by nonlinearity even if X_{3E} is empty.

Moreover, as demonstrated in Proposition S1 of the online Supplement, consistency of PARA depends on whether or not X_1 is present in (1). When X_1 is non-empty, which is the most common setting in practice, consistency generally fails due to the correlation between X_1 and the error in the linear projection of X_2 on X_3 . This result, which we believe is new, further favors PILS, whose consistency does not depend on X_1 being empty.

2.5.3 Comparison with MSII(-FM)

While PILS and MSII(-FM) are both two-sample semiparametric estimators, the former has three novel features relative to the latter. First, while both MSII(-FM) and PILS are \sqrt{n} -consistent and asymptotically normal two-sample estimators, their approaches to restoring consistency differ. Consistency of MSII(-FM) is established by imputing X_2 from \mathcal{S}_2 and

then eliminating the non-vanishing bias caused by the imputation. The bias correction requires a consistent, nonparametric estimate of Σ_2 , and the estimation error from $\hat{\Sigma}_2$ is $O_p(n^{-1/2})$. As a consequence, the asymptotic variance of MSII(-FM) tends to be large and highly complicated because of multiple asymptotically normal terms with the same $O_p(n^{-1/2})$ rate.

In contrast, consistency of PILS comes from directly imputing a consistent estimate of $g_2(X_3) = E(X_2|X_3)$ in place of the missing X_2 . By construction, it does not require bias correction that is a key ingredient in MSII(-FM).² The asymptotic variance of PILS also comes from two $O_p(n^{-1/2})$ terms, namely, the sampling error as in OLS and the approximation error from replacing the unobservable g_2 with its kernel estimate \hat{g}_2 . Although it is hard to compare asymptotic variances of PILS and MSII(-FM) analytically, Monte Carlo simulations below indicate that the former tends to be smaller than the latter.

Second, the asymptotic analysis we develop for PILS explicitly incorporates discrete matching variables. This is in contrast to the asymptotic analysis of Hirukawa and Prokhorov (2018), who do not accommodate discrete variables explicitly but argue that, similarly to the treatment effect literature (see, e.g., Abadie and Imbens, 2006), the inclusion of discrete matching variables with a finite number of support points does not affect convergence rates of MSII(-FM).

Third, we clarify the role of excluded continuous matching variables for identification. Hirukawa and Prokhorov (2018) maintain the assumption that all common variables enter the regression and are used for both estimation and matching. The price to pay for this assumption is that we need to impose nonlinearity of the conditional mean of the missing regressor given the matching variables in order to achieve identification. This paper relaxes the assumption so that some common variables may be employed only for imputation. The existence of such common variables allows for linearity in the conditional mean.

²From the viewpoint of imputation, the relationship between MSII(-FM) and PILS may be compared to the one between the NNM-based estimators studied by Abadie and Imbens (2006, 2011) and kernel-based matching estimators studied by Heckman, Ichimura and Todd (1998).

3 Finite-Sample Performance

3.1 Monte Carlo Setup

In this section we present selected results of an extensive Monte Carlo study comparing finite-sample properties of PILS with the alternative estimators. The design of the Monte Carlo study is meant to reflect the two scenarios depending on whether or not X_1 is empty. Due to the limitation of space, we focus on the more realistic case in which X_1 is non-empty. Monte Carlo designs and results when X_1 is empty are available in the online Supplement.

The Monte Carlo study is designed to mimic important aspects of the return to schooling application in Section 4. These are: (a) allowing for both included and excluded continuous matching variables (*educ* and *age* in the application); (b) maintaining $n/m = 2$ and considering the sample sizes of $(n, m) = (2000, 1000)$; (c) permitting the included continuous matching variable X_{3IC} (*educ*) to become endogenous when X_2 (*abil*) is omitted; and (d) setting the correlation between the endogenous variable X_{3IC} and the instrument Z (*fatheduc*) equal to 0.4. Of primary interest is estimation of β_3 , the coefficient on X_{3IC} (*educ*).

In this Monte Carlo study, the data are generated as follows. Consider the linear regression model

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_{3IC} + u, \quad (2)$$

where true parameter values are $\beta_0 = \beta_1 = \beta_2 = \beta_3 = 1$, and β_3 is the parameter of interest. Reduced forms of the regressors X_1 and X_2 are

$$\begin{aligned} X_1 &= 1 + X_{3IC}^2 + X_{3EC} + 2X_{3ED} + \eta_1, \text{ and} \\ X_2 &= h(X_{3IC}) - \frac{3}{4}(H_2 + 1)X_{3EC} + 2X_{3ED} + \eta_2 \end{aligned}$$

for some function $h(\cdot)$ to be specified shortly, where $H_r := E\{X_{3IC}^r h(X_{3IC})\}$ for $r \geq 0$, and X_{3IC} , X_{3EC} and X_{3ED} are included continuous, excluded continuous and excluded discrete common variables, respectively. The regressor X_2 is missing in the primary sample \mathcal{S}_1 . In this scenario, if we omit X_2 from (2) then we estimate the short regression

$$Y = \beta_0 + \beta_1 X_1 + \beta_3 X_{3IC} + (u + \beta_2 X_2) =: \beta_0 + \beta_1 X_1 + \beta_3 X_{3IC} + v. \quad (3)$$

Because X_{3IC} is correlated with X_2 , we generate the instrument for X_{3IC} by $Z = \pi (X_{3IC} - 2H_1X_{3ED}) + \xi$, where $\pi = \sqrt{3\rho^2 / \{4 - \rho^2 (4 + 3H_1^2)\}}$ for $\rho = \text{corr}(X_{3IC}, Z) \in \{0.4, 0.1\}$. Notice that $\rho = 0.4$ and 0.1 correspond to the cases of strong and weak instruments, respectively. We also make the following distributional assumptions:

$$\begin{aligned} (u, \eta_1, \eta_2, \xi)' &\sim N(0_{4 \times 1}, I_4); \\ X_{3IC}, X_{3EC} &\stackrel{iid}{\sim} U[-2, 2]; \\ X_{3ED} &\sim \text{Bernoulli}(1/2) - 1/2 = \begin{cases} -1/2 & \text{w.p. } 1/2 \\ 1/2 & \text{w.p. } 1/2 \end{cases}; \text{ and} \\ u \perp X_3 &:= (X_{3IC}, X_{3EC}, X_{3ED}). \end{aligned}$$

The function $h(\cdot)$ is specified as one of the following:

$$h(x) = \begin{cases} x & \text{[Model A]} \\ \sin(\pi x/4) & \text{[Model B]} \\ \{(5/3)x + 3/2\} \mathbf{1}\{x < 0\} + \{-(4/3)x + 3/2\} \mathbf{1}\{x \geq 0\} & \text{[Model C]} \\ x + (5/\tau)\phi(x/\tau) - (5/2)\{\Phi(2/\tau) - 1/2\}, \tau = 3/4 & \text{[Model D]} \end{cases},$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the pdf and cdf of $N(0, 1)$, respectively. Models A and B are monotone. Linearity of $h(\cdot)$ in Model A does not cause identification failure, because two common variables (X_{3EC}, X_{3ED}) are excluded from (2). Model B is nonlinear and odd. Models C and D are non-monotone, and the former is even non-smooth when $x = 0$. Inspired by the Monte Carlo design of Horowitz and Spokoiny (2001), Model D can be viewed as a linear function with a bump. It follows from the distributional assumption on X_{3IC} that values of (H_1, H_2, H_3) are:

$$(H_1, H_2, H_3) = \begin{cases} (0, 4/3, 0) & \text{[Model A]} \\ (0, 8/\pi^2, 0) & \text{[Model B]} \\ (0, 2/9, -1) & \text{[Model C]} \\ (0, 4/3, -0.9989) & \text{[Model D]} \end{cases}.$$

The above procedure provides us with two observable samples

$$\mathcal{S}_1 = \{(Y_i, X_{1i}, X_{3ICi}, X_{3ECi}, X_{3EDi}, Z_i)\}_{i=1}^n \text{ and } \mathcal{S}_2 = \{(X_{2j}, X_{3ICj}, X_{3ECj}, X_{3EDj})\}_{j=1}^m.$$

The complete sample

$$\mathcal{S}^* = \{(Y_i, X_{1i}, X_{2i}, X_{3ICi}, X_{3ECi}, X_{3EDi}, Z_i)\}_{i=1}^n$$

is the sample that would not be observed in practice. Finally, the matched sample

$$\mathcal{S} = \{(Y_i, X_{1i}, X_{2j_1(i)}, \dots, X_{2j_K(i)}, X_{3ICi}, X_{3ECi}, X_{3EDi}, Z_i)\}_{i=1}^n$$

is constructed via the NNM with respect to $X_3 = (X_{3IC}, X_{3EC}, X_{3ED})$ as in Hirukawa and Prokhorov (2018). The NNM is based on the Mahalanobis distance, and the number of matches is set equal to $K = 1$ (single match) that is most commonly chosen. Sample sizes of \mathcal{S}_1 and \mathcal{S}_2 are $n \in \{500, 1000, 2000\}$ and $m = n/2$. For each combination of the functional form of $h(\cdot)$ and value of π , we draw 1000 Monte Carlo samples.

There are three options to estimate β_3 consistently. The first option is to estimate the short regression (3) using \mathcal{S}_1 only. While $E(v) = E(X_1v) = 0$ holds, $E(X_{3IC}v) \neq 0$ (i.e., X_{3IC} is endogenous) is the case and OLS for (3) is inconsistent. However, it turns out that $\rho \neq 0$ and $E(Zv) = 0$. Hence, we can estimate β_3 consistently by using Z as an instrument for X_{3IC} . The remaining two options rely on both \mathcal{S}_1 and \mathcal{S}_2 to estimate the long regression in (2). In the second option, after constructing the matched sample \mathcal{S} from \mathcal{S}_1 and \mathcal{S}_2 , we run MSII-FM using \mathcal{S} . As the third option we employ PILS using both \mathcal{S}_1 and \mathcal{S}_2 .

Based on the above estimation strategies, we compare the following estimators of β_3 : (i) the infeasible OLS estimator for (2) using \mathcal{S}^* [OLS*];³ (ii) the inconsistent OLS estimator for (3) using \mathcal{S}_1 only [OLS-S]; (iii) the IV estimator for (3) using \mathcal{S}_1 only [IV-S]; (iv) the inconsistent MSOLS estimator for (2) using \mathcal{S} [MSOLS]; (v) the MSII-FM estimator for (2) using \mathcal{S} and second, third and fourth-order polynomials, where the first-step MSII estimator is reported as a benchmark [MSII, MSII-FM]; (vi) the PARA estimator for (2) using \mathcal{S}_1 and \mathcal{S}_2 [PARA]; (vii) the PILS estimator for (2) using \mathcal{S}_1 , \mathcal{S}_2 and the Epanechnikov kernel $K_E(t) = (3/4)(1 - t^2) \mathbf{1}\{|t| \leq 1\}$ [PILS-E];⁴ and (viii) the PILS estimator for (2) using \mathcal{S}_1 , \mathcal{S}_2 and the beta kernel [PILS-B].

A few additional remarks are in order. First, implementing the two PILS estimators requires a choice of the smoothing parameters. We put $\hat{h} = \hat{\sigma}_X (\log m/m)^{0.3}$ for the Epanechnikov kernel, $\hat{b} = \hat{\sigma}_U (\log m/m)^{0.6}$ for the beta kernel, and $\hat{\lambda} = (\log m/m)^{0.6}$ for the discrete kernel, where $\hat{\sigma}_X$ and $\hat{\sigma}_U$ are sample standard deviations of the continuous common variable

³OLS* is consistent, because $E(u) = E(X_1u) = E(X_2u) = E(X_{3IC}u) = 0$.

⁴In addition to the second-order Epanechnikov kernel, we consider the fourth and sixth-order Epanechnikov kernels by Hansen (2005), where their functional forms are

$$K_{E,4}(t) = \frac{15}{8} \left(1 - \frac{7}{3}t^2\right) K_E(t) \text{ and } K_{E,6}(t) = \frac{175}{64} \left(1 - 6t^2 + \frac{33}{5}t^4\right) K_E(t),$$

respectively. However, our early simulation results indicate huge finite-sample biases from these estimators, and thus are not reported.

in the original scale $X (= X_{3IC}, X_{3EC})$ and in the transformed scale $U := (X + 2) / 4 \in [0, 1]$, respectively. These shrinkage rates indeed fulfill the requirements in Section 2.4.4.

Second, the functional form of the imputed regressor in PARA matches the reduced form of X_2 for Model A only. However, even if the reduced form of X_2 is misspecified, it is still possible to estimate regression (2) consistently by PARA. It turns out that under our Monte Carlo design, PARA is consistent if and only if $h(\cdot)$ is odd, i.e., for both Models A and B; see Appendix A.5 for a proof of this claim, which is of independent interest.

For each estimator of β_3 , the following performance measures are computed: (i) *Mean* (simulation average of the parameter estimate); (ii) *SD* (simulation standard deviation of the parameter estimate); (iii) *RMSE* (root mean-squared error of the parameter estimate); (iv) \overline{SE} (simulation average of the standard error); and (v) *CR* (coverage rate for the nominal 95% confidence interval). Formulae for standard errors are as follows: (a) heteroskedasticity-robust standard errors by Eicker (1963) and White (1980) are calculated for OLS* and IV-S; (b) the one for MSII-FM follows Theorem 4 of Hirukawa and Prokhorov (2018); (c) the one for PILS appears in Theorem 2; and (d) the one for PARA is given in Proposition S2 of the online Supplement. Because of inconsistency we do not compute \overline{SE} or *CR* of OLS-S and MSOLS. Those of the first-step MSII are also omitted due to its nonparametric convergence rate. Likewise, we present \overline{SE} and *CR* of PARA only when it becomes consistent, i.e., for Models A and B.

TABLE 1 ABOUT HERE

3.2 Results

Simulation results are summarized in Table 1. To save space, we present the results for $(n, m) = (1000, 500), (2000, 1000)$; see the online Supplement for comprehensive results. In addition, there is not much difference in the results for different orders of the polynomial in MSII-FM, and thus only the results from a second-order polynomial are reported.

OLS* is an infeasible, oracle estimator, and we should think of it merely as a benchmark. It is of more practical importance and interest to compare efficiency among other consistent estimators of β_3 , which are all less efficient than OLS*. However, the reasons for their efficiency losses differ between IV-S and two-sample estimators including MSII-FM, PILS

and PARA (for Models A and B only). The efficiency loss of IV-S is due to leaving out the missing regressor X_2 and relying on an instrument, whereas that of the two-sample estimators is attributed to imputing X_2 from another data source. Moreover, IV-S using a valid but weak instrument ($\rho = 0.1$) is severely biased and has huge standard deviations (and standard errors) even when $(n, m) = (2000, 1000)$. When the available instrument is weak, the advantage of two-sample estimators is more pronounced.

Now we turn to the two-sample estimators. MSII-FM appears to be less efficient than PILS. Furthermore, PILS-B is consistently centered around the true value, whereas the finite-sample bias of PILS-E is often substantial. Therefore, it appears that PILS-B is a safer choice than PILS-E. PARA may look less biased and more efficient than MSII-FM and PILS. However, this occurs only when PARA is consistent (i.e., for Models A and B), and only by a slight margin compared with PILS-B. For other models, it generates a considerable bias.

All in all, PILS-B appears to be an efficient estimation strategy within the regression framework with missing regressors. Although its standard deviations are always greater than those of OLS*, the relative efficiency loss can be thought of as the price to pay for identifying and estimating the regression using two samples jointly. It may be tempting to rely on PARA, in particular, when X_1 is empty. This is the best case scenario for PARA, because in this case it is always consistent. Nonetheless, as the online Supplement suggests, even in this case, PILS-B outperforms PARA in terms of efficiency (SD) and quality ($RMSE$). Therefore, it is safe to employ PILS-B whenever two samples are available, regardless of whether X_1 is empty or not.⁵

4 An Application: Estimation of Return to Schooling

4.1 Earnings Function

Estimation of the causal link between education and earnings has been a focus of labor economists over the last several decades. Card (2001) suggested that endogeneity of education in the earnings equation might at least in part be responsible for the continuing interest

⁵In the online Supplement, we report how PILS compares with nonlinear IV of a kind similar to Dieterle and Snell (2016). Again, PILS-B performs quite well in this setup.

in uncovering the causal effect of education on labor market outcomes. As Griliches (1977) discusses in detail, empirical labor economists have long speculated that the primary reason why education is endogenous when estimating the returns to schooling is some omitted explanatory variable(s) – unobservable factors that influence education such as ability and motivation – that are likely to have a direct effect on individual earnings and wages. The aim of this section is to estimate the rate of return to additional schooling using alternative estimators available to us. These include the traditional one-sample OLS and linear and nonlinear IV estimators as well as MSII, PARA, and PILS. A more comprehensive comparison can be found in the online Supplement.

Following the classical framework of the human capital earnings/wage function of Mincer (1974) we assume additivity and wish to use US data in order to estimate the causal effect of education on earnings using the following model:

$$\begin{aligned} \log(\textit{earnings}) = & \beta_0 + \beta_1\textit{education} + \beta_2\textit{ability} + \beta_3\textit{experience} + \beta_4\textit{experience}^2 \\ & + \beta_5\textit{married} + \beta_6\textit{black} + \beta_7\textit{south} + \beta_8\textit{urban} + u, \end{aligned} \quad (4)$$

where $\log(\textit{earnings})$ is the natural logarithm of the individual’s total annual labor income, $\textit{education}$ is the person’s completed years of education, $\textit{ability}$ is the individual’s ability/skills (with zero mean), $\textit{experience}$ is work experience of the person, $\textit{married}$ is an indicator for whether the individual is married, \textit{black} is an indicator for whether the person is black, \textit{south} is an indicator for whether the person currently lives in the southern geographical region, \textit{urban} is an indicator of the individual’s urban residence while growing up,⁶ and u is an idiosyncratic error.

The main issue with estimating equation (4) is that econometricians are typically unable to observe the individual’s skills. Because of that, the (one-sample) OLS estimator of the return to schooling – β_1 – from equation (4) where $\textit{ability}$ is excluded is likely to suffer from an “ability bias.” Two textbook solutions to this problem are (1) to find within the same data set a valid proxy for the unobservable skills and use OLS estimation, and (2) to find within the same data set a valid instrumental variable for the individual’s educational level and use the IV approach. Since in addition to these two approaches, we advocate

⁶The publicly available part of the PSID survey that we use does not provide information on current urban residence.

using two-sample estimation, we would like to be able to apply in practice the one- and two-sample approaches to the same applied task.

4.2 Constructing the Samples

To estimate earnings and wage equations for the US, labor economists frequently use such micro-data sets as CPS, PSID and the National Longitudinal Survey (NLS). While CPS has a larger sample and is more representative of the entire demographic composition of the US population than the other two surveys, it does not usually collect such important wage determinants as actual work experience and ability. PSID does provide information on actual work experience but, similar to CPS, generally does not collect data on ability. NLS routinely reports ability measures and contains data on actual and potential work experience but it is less representative of the US labor force. Given the features of these surveys, PSID and NLS seem most suitable as our main and auxiliary data sets, respectively.

We employ the 1972 wave of PSID as our main data set. We choose this wave so that we can estimate equation (4) using not only IV and PILS but also using a proxy for unobserved skills and ability. While PSID generally does not contain any measures of unobserved ability, the 1972 wave is among the few PSID waves that do include an ability measure.

In the 1972 wave, the PSID respondents were administered a particular assessment of abstract thinking – the Lorge-Thorndike Intelligence Test (LTIT). In essence, this is a test of verbal skills – a sentence completion test – that Veroff et al. (1971, p.26) describe “as a feasible, reasonably valid assessment of what psychologists have labeled intelligence.” Furthermore, Veroff et al. (1971, p.26) advise that this test “seems to correlate well with most different kinds of tests of intelligence, well enough to suggest using it singly without going to multiple measurement.” LTIT administered to the PSID respondents contained 13 questions, where each question received a point for the correct response. Thus, the score for the entire sentence completion test can range from zero (the worst outcome) to 13 (the best outcome). We call the variable containing the total test score *IQ score*, and we demean it.

In addition to the information on the individual’s ability discussed above, we also gather information on the individual’s actual work experience, completed years of education, age,

whether the person is black, married, and lives in the southern geographical region⁷. Furthermore, we include a dummy variable for whether the person grew up in an urban area and for whether the person grew up in the southern region. Finally, in our PSID sample, we also get the information on the completed years of education by the individual’s father. The last variable – father’s educational level – is used as an instrument for the individual’s educational attainment when no measure of ability is available.

TABLE 2 ABOUT HERE

Our main sample contains 2430 men who reported positive labor income in 1971. Panel A of Table 2 reports summary statistics for the variables in our PSID sample. The sample correlation between an individual’s education and the father’s education is 0.434.

To exploit the two-sample approaches, we need a second sample where the ability measure is available and/or potentially more reliable. We employ the CARD data set provided with Wooldridge (2013). This data set contains observations from the 1977 wave of the Young Men Cohort of NLS. We exploit observations with positive wages in 1976.

As an ability measure we employ the results of the “Knowledge of the World of Work” (*KWW score*) test that the NLS respondents were administered during their interviews in 1966. The KWW score from NLS is arguably a better measure of unobserved ability than the IQ score from the 1972 wave of PSID. A higher KWW score indicates higher intelligence. We demean this measure before using it to estimate equation (4).

Panel B of Table 2 provides summary statistics for the NLS sample of $m = 1102$ respondents when *education*, *married*, *black*, *south*, and *urban* are used as the included common variables, X_{3I} , and *age* (the individual’s age, in years) and *south while growing up* (an indicator for whether the person resided in the southern region while growing up) are used as the excluded common variables, X_{3E} . Note that *experience* and *experience*², which are not observed in the NLS sample, represent X_1 in regression (4). The summary statistics for *experience* and *experience*² are provided in Panel A of Table 2.

⁷We follow the 1979 NLS of Youth definition of the southern region. The southern region includes Alabama, Arkansas, Delaware, District of Columbia, Florida, Georgia, Kentucky, Louisiana, Maryland, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia, and West Virginia.

4.3 Estimation Approaches

We report estimation results in Table 3.⁸ Columns (1)–(3) report the parametric estimation results based on just the PSID sample. Column (1) reports the OLS estimates using the ability proxy from the 1972 wave. Column (2) reports the OLS results that do not account for unobserved ability. In most real-life settings, the results in column (1) are infeasible and those in column (2) are subject to the “ability bias.” Column (3) provides the IV estimates where the father’s education is used as an IV for an individual’s educational level.

Columns (4)–(7) report two-sample estimates based on both the PSID and NLS samples. Columns (4), (5), and (7) contain the results based on semiparametric two-sample approaches, while column (6) reports the only estimates based on a fully parametric two-sample procedure. Specifically, column (4) contains the OLS estimates when the two samples are combined using hot-deck imputation and no bias correction is applied. As discussed by Hirukawa and Prokhorov (2018), this estimator is inconsistent. Column (5) provides the MSII-FM results that account for the imputation biases. In essence, the MSII approach is a bias-corrected version of the MSOLS approach. Column (6) exhibits the PARA results as yet another two-sample estimator. It is worth emphasizing that by Proposition A2, PARA is not guaranteed to be consistent in this case. Finally, column (7) reports the PILS results using the beta kernel, chosen on the basis of the simulation results in previous section.

The two-sample estimators use five variables common to both samples – *education*, *married*, *black*, *south*, and *urban* – as included variables and two variables – *age* and *south while growing up* – as excluded variables. Note that since in our case both the main (PSID) and auxiliary (NLS) samples contain information on the individual’s educational level, we treat *education* as a part of X_3 . We choose not to employ *experience* as a common variable because experience in PSID represents actual experience, while experience in NLS represents potential experience, which is quite different. As a consequence, our entire set of common variables contains *education*, *married*, *black*, *south*, *urban*, *age*, *south while growing up*, where *education* and *age* are treated as continuous.

The NNM used by MSOLS and MSII-FM adopts a single match ($K = 1$) based on the

⁸GAUSS codes implementing the IV and two-sample estimators as well as the data sets are available at <https://sites.google.com/site/artemprokhorov0/papers/SimulationAndApplicationFiles.zip>.

Mahalanobis metric. We average the (demeaned) KWW score for ties in our second (NLS) sample and assign this average as a unique value of the ability measure to the respondent with the given values of our common variables. As a consequence, $m = 1102$ respondents remain in our \mathcal{S}_2 . A second-order polynomial is used in MSII-FM.

Our choice of a kernel for PILS reflects the favorable performance of the beta kernel in simulations. For PILS, each continuous common variable X is converted from its original scale to a variable $U := (X - m_X) / (M_X - m_X) \in [0, 1]$, where M_X and m_X are maximal and minimal values in the pooled sample constructed from observations of X in \mathcal{S}_1 and \mathcal{S}_2 . The reason for using the maximum and minimum of the pooled sample is that the ranges of *education* and *age* are quite different between \mathcal{S}_1 and \mathcal{S}_2 . We use $\hat{b} = \hat{\sigma}_U (\log m/m)^{0.6}$ and $\hat{\lambda} = (\log m/m)^{0.6}$ for the beta and discrete kernels, respectively, where $\hat{\sigma}_U$ is the sample standard deviation of a converted continuous common variable U in \mathcal{S}_2 .

4.4 Empirical Findings

It is immediately clear from Table 3 that the PARA estimation results substantially differ from the rest. First, contrary to the corresponding OLS* estimate, the estimate of the rate of return to education is negative. Second, the PARA estimate of the coefficient on *ability* is expectedly positive but extraordinarily large when compared to the OLS* estimate. Third, PARA estimates are quite unstable in that the standard error of each coefficient estimate is much larger than those of the corresponding estimates from other methods in general. Because it is hard to provide a meaningful interpretation to the PARA results, we exclude them from further discussion. One lesson from the PARA estimates is that parametric imputation methods for missing regressors should be used with caution.

Table 3 shows that the signs of coefficient estimates on all the regressors except *ability* are as expected and most of the estimates are significant. The MSOLS estimator yields negative and insignificant estimates of the ability effect, whereas the MSII estimator results in a positive and insignificant estimate. Interestingly, OLS* and PILS are the only two estimators that produce positive (as one would expect) and statistically significant estimates of the ability effect. Finally, we note that the PILS standard errors tend to be smaller than those of MSII, as predicted in a previous section.

TABLE 3 ABOUT HERE

Next we focus on the estimates of the rate of return to education. The estimates reported in Table 3 provide evidence that the father’s education is unlikely to be a valid instrument for an individual’s education. This is because if the father’s education were a valid IV, the OLS-S and MSOLS estimators of β_1 would be upward-inconsistent. However, we observe that the IV estimate is the largest in magnitude. Furthermore, if this instrument is valid and there is no measurement error in the PSID ability measure then OLS* and IV are both consistent. However, we see that the OLS* and IV estimates are noticeably different, suggesting again that the two estimators are unlikely to be both consistent.

The infeasible OLS – OLS* – is the second smallest (conceding in magnitude only to PILS). In addition, if there is no measurement error in the PSID ability measure then OLS* is consistent, but OLS-S and MSOLS are still upward-inconsistent. This possibility seems to fit well with our estimates in Table 3. Indeed, the feasible one-sample estimate of β_1 with omitted ability – OLS-S – is between IV and infeasible OLS. In fact, the MSOLS and MSII-FM estimates of β_1 are also between the one-sample IV and infeasible OLS* estimates. Therefore, there is ground to view the infeasible OLS* as a benchmark result in our analysis.

Importantly, the only estimate of β_1 that is smaller than the infeasible OLS is PILS. We note that PILS is also the second closest (in absolute value) to the infeasible OLS estimate, while the first closest is MSII-FM. This would be expected given that only MSII-FM and PILS are the consistent approaches if the instrument is invalid. Further note, that the possibility that PILS (and not OLS*) is closer to the true return to education is not unrealistic. If the PSID ability measure is error-ridden, OLS* is upward-inconsistent suggesting that PILS is likely to be the only estimator that delivers the closest estimate to the true population value.⁹

5 Concluding Remarks

When some regressors are unavailable for a regression analysis, econometricians often resort to IV estimation. They typically make considerable effort to find and justify valid instru-

⁹The online Supplement presents a detailed discussion on this claim.

ments for the regressors that are suspected to be endogenous due to their possible correlation with the omitted regressors. In this paper we have explored two-sample alternatives to this conventional approach.

We have developed the PILS estimation procedure for such models. The procedure first uses an auxiliary sample to obtain a nonparametric estimator of the conditional mean of the regressor which is missing in the main sample. Variables that are common in both samples are used in the conditioning set. In the second step of the procedure, we simply run OLS using a plug-in estimates of the conditional mean.

We establish asymptotic normality of PILS. Attractive finite-sample properties of PILS are confirmed in a Monte Carlo study. In particular, simulations demonstrate numerically that PILS is generally more efficient than IV and it dominates other available estimators, e.g., MSII(-FM), in terms of mean squared error. The results also indicate superior finite-sample properties of PILS using the beta kernel.

However, in order for PILS to attain the parametric rate of convergence, the number of continuous common variables must be three or less (provided that nonnegative kernels are employed). In this respect, the curse of dimensionality plays out similarly to the NNM-based estimation considered by Hirukawa and Prokhorov (2018).

A natural extension is to adopt a dimension reduction method when using multiple common variables. This can be done using a sparsity-based algorithm, e.g., lasso of the first step, or using a propensity score algorithm. We leave these ideas for future research.

A Appendix: Technical Proofs

A.1 Construction of Product Kernel for Estimating $g_2(\cdot)$

We employ a continuous univariate kernel for each of d_{3C} continuous variables $X_{3C,p}$, $p = 1, \dots, d_{3C}$. For simplicity, suppose that $X_{3C,p}$ is smoothed by a univariate symmetric kernel $K(\cdot)$ and bandwidth h_p . Then, the product kernel for $X_{3C,p}$ is

$$\mathcal{K}(t_{3C}; x_{3C}, \mathbf{h}) = \prod_{p=1}^{d_{3C}} \frac{1}{h_p} K\left(\frac{t_{3C,p} - x_{3C,p}}{h_p}\right),$$

where $t_{3C} := (t_{3C,1}, \dots, t_{3C,d_{3C}})$, $x_{3C} := (x_{3C,1}, \dots, x_{3C,d_{3C}})$ and $\mathbf{h} := (h_1, \dots, h_{d_{3C}})$ are vectors of data points, design points and bandwidths, respectively.

Next, we construct a kernel for the discrete component. While a variety of discrete kernels can be applied for this component, our focus is on those given by Aitchison and Aitken (1976). In what follows a product of their discrete kernels is exclusively considered. Each of d_{3D} discrete variables $X_{3D,q}$, $q = 1, \dots, d_{3D}$ is assumed to take $r_q (\geq 2)$ different values, i.e., $X_{3D,q} \in \{0, 1, \dots, r_q - 1\}$. In addition, each discrete variable is classified into either unordered or ordered, because the kernels employed for the two types of categorical variables differ slightly. The univariate discrete kernel for an *unordered* variable is

$$l(t_{3D,q}; x_{3D,q}, \lambda_q) := \begin{cases} 1 & \text{if } t_{3D,q} = x_{3D,q} \\ \lambda_q / (r_q - 1) & \text{if } t_{3D,q} \neq x_{3D,q} \end{cases},$$

where $t_{3D,q}$, $x_{3D,q}$ and $\lambda_q \in (0, 1)$ are the data point, design point and bandwidth, respectively. The univariate discrete kernel for an *ordered* variable takes the form of

$$\ell(t_{3D,q}; x_{3D,q}, \lambda_q) := \binom{r_q}{|t_{3D,q} - x_{3D,q}|} (1 - \lambda_q)^{r_q - |t_{3D,q} - x_{3D,q}|} \lambda_q^{|t_{3D,q} - x_{3D,q}|}.$$

If there are $q_1 (\leq d_{3D})$ unordered discrete variables, then the product kernel for all d_{3D} discrete variables is given by

$$\mathcal{L}(t_{3D}; x_{3D}, \lambda) = \left\{ \prod_{q=1}^{q_1} l(t_{3D,q}; x_{3D,q}, \lambda_q) \right\} \left\{ \prod_{q=q_1+1}^{d_{3D}} \ell(t_{3D,q}; x_{3D,q}, \lambda_q) \right\},$$

where $t_{3D} := (t_{3D,1}, \dots, t_{3D,d_{3D}})$, $x_{3D} := (x_{3D,1}, \dots, x_{3D,d_{3D}})$ and $\lambda := (\lambda_1, \dots, \lambda_{d_{3D}})$.

For the continuous component, we also apply an alternative nonstandard kernel. Taking compactness of \mathbb{X}_{3C} into account (see Assumption 2), we employ the beta kernel by Chen (1999) in place of the univariate symmetric kernel. For a design point $x \in [0, 1]$ and the smoothing parameter b , the beta kernel is defined as

$$K_{B(x,b)}(t) = \frac{t^{x/b} (1-t)^{(1-x)/b}}{B\{x/b+1, (1-x)/b+1\}} \mathbf{1}\{t \in [0, 1]\}.$$

It turns out that convergence properties of nonparametric estimators based on this kernel family have not been established. In a companion paper (Hirukawa, Murtazashvili and Prokhorov, 2020), we provide a careful proof of weak and strong uniform convergence with rates of nonparametric estimators smoothed by the beta kernel.

A.2 Proof of Theorem 1

The proof requires a lemma on uniform consistency of the kernel regression estimator $\hat{g}_2(\cdot)$. Let $g_{2k}(\cdot)$ be the k th element of $g_2(\cdot)$ for $k = 1, \dots, d_2$. In order to differentiate the continuous kernels used in the nonparametric estimation of $g_{2k}(\cdot)$, we denote the estimates obtained using ν th-order symmetric and beta kernels by $\hat{g}_{2k}^{S\nu}(\cdot)$ and $\hat{g}_{2k}^B(\cdot)$, respectively. Then, the lemma below establishes weak uniform consistency with rates of $\hat{g}_{2k}^{S\nu}(\cdot)$ and $\hat{g}_{2k}^B(\cdot)$.

Lemma A1. *If Assumptions 1–6 hold, then, for $k = 1, \dots, d_2$, as $m \rightarrow \infty$,*

$$\begin{aligned} \sup_{x_3 \in \mathbb{X}_3} |E\{\hat{g}_{2k}^{S\nu}(x_3)\} - g_{2k}(x_3)| &= O\left(\sum_{p=1}^{d_{3C}} h_p^\nu + \sum_{q=1}^{d_{3D}} \lambda_q\right), \\ \sup_{x_3 \in \mathbb{X}_3} |\hat{g}_{2k}^{S\nu}(x_3) - E\{\hat{g}_{2k}^{S\nu}(x_3)\}| &= O_p\left(\sqrt{\frac{\log m}{m \prod_{p=1}^{d_{3C}} h_p}}\right), \\ \sup_{x_3 \in \mathbb{S}_{X_{3C}} \times \mathbb{X}_{3D}} |E\{\hat{g}_{2k}^B(x_3)\} - g_{2k}(x_3)| &= O\left(\sum_{p=1}^{d_{3C}} b_p + \sum_{q=1}^{d_{3D}} \lambda_q\right), \text{ and} \\ \sup_{x_3 \in \mathbb{S}_{X_{3C}} \times \mathbb{X}_{3D}} |\hat{g}_{2k}^B(x_3) - E\{\hat{g}_{2k}^B(x_3)\}| &= O_p\left(\sqrt{\frac{\log m}{m \sqrt{\prod_{p=1}^{d_{3C}} b_p \theta_p}}}\right). \end{aligned}$$

A.2.1 Proof of Lemma A1

First two statements on $\hat{g}_{2k}^{S\nu}(\cdot)$ are implied by Theorem 2.1 of Li and Ouyang (2005). Remaining two statements on $\hat{g}_{2k}^B(\cdot)$ are established by Theorem 6 of Hirukawa, Murtazashvili and Prokhorov (2020) with $r_n = O(1)$ due to uniform boundedness from below of $f(\cdot)$. ■

A.2.2 Proof of Theorem 1

To save space, we concentrate only on the case with the beta estimator $\hat{g}_2^B(\cdot)$. The PILS estimator admits the expansion

$$\begin{aligned} \hat{\beta}_{PILS} &= \beta + S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, \epsilon} + S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, \{g_2(X_3) - \hat{g}_2^B(X_3)\}'} \beta_2 \\ &= \beta + S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, [g_2(X_3) - E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2 \\ &\quad + S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, \epsilon} - S_{n, \hat{X}, \hat{X}}^{-1} S_{n, \hat{X}, [\hat{g}_2^B(X_3) - E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2, \end{aligned} \tag{A1}$$

where the third and fourth terms on the right-hand side can be interpreted as those of the sampling error in regression (1) and the approximation error of g_2 , respectively. It follows from Lemma A1 that $B_{g_2} := S_{n,\hat{X},\hat{X}}^{-1} S_{n,\hat{X},[g_2(X_3)-E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2 = O_p\left(\sum_{p=1}^{d_{3C}} b_p + \sum_{q=1}^{d_{3D}} \lambda_q\right)$. Each of the last two terms is $O_p(n^{-1/2})$ by the central limit theorem (CLT); see the proof of Theorem 2 below for more details. Therefore,

$$\hat{\beta}_{PILS} = \beta + O_p\left(\sum_{p=1}^{d_{3C}} b_p + \sum_{q=1}^{d_{3D}} \lambda_q\right) + O_p(n^{-1/2}) + O_p(n^{-1/2}) \xrightarrow{p} \beta$$

as $n, m \rightarrow \infty$. ■

A.3 Proof of Theorem 2

We continue to focus on the case with the beta estimator $\hat{g}_2^B(\cdot)$. Short-hand notations such as $g_{2i} = g_2(X_{3i})$, $E_i(\hat{g}_{2i}^B) = E\{\hat{g}_2^B(X_{3i})|X_{3i}\}$, and $\mathcal{W}_{j,i} = \mathcal{W}(X_{3j}; X_{3i}, \mathbf{b}, \lambda)$ also apply. Then, by Lemma A1 and CLT,

$$S_{n,\hat{X},\hat{X}} = S_{n,X,X} + o_p(1) \xrightarrow{p} \Phi_{X,X} > 0, \text{ and} \quad (\text{A2})$$

$$\sqrt{n}S_{n,\hat{X},\epsilon} = \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i + o_p(1) \xrightarrow{d} N(0_{(d+1) \times 1}, \Omega_1) \quad (\text{A3})$$

for the sampling error part. Moreover,

$$S_{n,\hat{X},[\hat{g}_2^B(X_3)-E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2 = S_{n,X,[\hat{g}_2^B(X_3)-E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2 + S_{n,\hat{X}-X,[\hat{g}_2^B(X_3)-E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2$$

for the approximation error part, where the second term is of a smaller order of magnitude than the first term. Also observe that $\hat{g}_{2i}^B - E_i(\hat{g}_{2i}^B) = \{\tilde{g}_{2i}^B - E_i(\tilde{g}_{2i}^B)\} + (\tilde{\eta}_{2i}^B - \eta_{2i}) + \eta_{2i}$, where $\tilde{\varphi}_i^B := \sum_{j=1}^m \varphi_j \mathcal{W}_{j,i} / \sum_{j=1}^m \mathcal{W}_{j,i}$ for $\varphi \in \{g_2, \eta_2\}$. Then, by a similar argument to Proposition 9 of Stengos and Yan (2001) and CLT,

$$\sqrt{n}S_{n,\hat{X},[\hat{g}_2^B(X_3)-E\{\hat{g}_2^B(X_3)|X_3\}]}' \beta_2 = \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \eta_{2i}' \beta_2 + o_p(1) \xrightarrow{d} N(0_{(d+1) \times 1}, \Omega_2). \quad (\text{A4})$$

Finally, it is straightforward to see that for every i and k ,

$$E\left[\left(\hat{X}_i \epsilon_i\right) \left\{\hat{X}_k (\hat{g}_{2k}^B - E_k(\hat{g}_{2k}^B))' \beta_2\right\}'\right] = 0_{(d+1) \times (d+1)}. \quad (\text{A5})$$

The result is established by substituting (A2)–(A5) into (A1) and rearranging it. ■

A.4 Proof of Corollary 1

Again to save space, we provide the proof for the case with the beta estimator $\hat{g}_2^B(\cdot)$. To control the convergence rate of θ , let $\theta \propto b^{\delta/d_{3C}}$ for a sufficiently small $\delta > 0$ so that $b/\theta \rightarrow 0$ can be established. Because $B_{g_2} = O_p(b) + o_p(n^{-1/2})$, $\sqrt{n}(\hat{\beta}_{PILS} - \beta - B_{g_2}) = \sqrt{n}(\hat{\beta}_{PILS} - \beta) + o_p(1)$ holds if

$$nb^2 = O\left\{\left(\frac{n}{m}\right)(\log m)^{4\alpha} m^{1-4\alpha}\right\} \rightarrow 0.$$

This is the case if one of the following conditions is true: (i) $n/m \rightarrow \kappa$ and $\alpha > 1/4$; (ii) $n/m \rightarrow 0$ and $\alpha > 1/4$; or (iii) $n/m \rightarrow \infty$, $n/m^{4\alpha} \rightarrow 0$ at a polynomial rate and $\alpha > 1/4$.

It also follows from Assumption 6'(b) that

$$\frac{\log m}{m(b\theta)^{d_{3C}/2}} = O\left\{\left(\frac{\log m}{m}\right)^{1-\alpha(d_{3C}+\delta)}\right\} \rightarrow 0$$

should be the case. This holds if $1-\alpha(d_{3C}+\delta) > 0$ or $\alpha < 1/(d_{3C}+\delta) < 1/d_{3C}$. Therefore, for $\hat{\beta}_{PILS}$ to be \sqrt{n} -consistency, α must satisfy $1/4 < \alpha < 1/d_{3C}$. This range is non-empty if $d_{3C} \leq 3$. ■

A.5 Proof of Claim about Consistency of PARA in Simulations

Let $\text{proj}(X_2|X_3) = A\mathcal{X}_3 = A_0 + A_1X_{3IC} + A_2X_{3EC} + A_3X_{3ED}$. Then, since X_{3IC} , X_{3EC} , X_{3ED} , and η_2 are zero-mean mutually independent random variables, the reduced form of X_2 implies that all the following statements are true:

1. $A_0 = 0$, $A_2 = -(3/4)(H_2 + 1)$, $A_3 = 2$;
2. A_1 satisfies $A_1X_{3IC} = \text{proj}\{h(X_{3IC})|X_{3IC}\}$; and
3. The projection error $e = X_2 - A\mathcal{X}_3 = h(X_{3IC}) - A_1X_{3IC} + \eta_2$.

By Proposition S1 of the online Supplement, PARA is consistent if and only if $E(X_1e) = 0$. However, it follows from the reduced form of X_1 and distributional assumptions on X_{3IC} , X_{3EC} , X_{3ED} , and η_1 that $E(X_1e) = H_2 = E\left\{X_{3IC}^2 h(X_{3IC})\right\}$ holds. The right-hand side is zero if and only if $h(\cdot)$ is odd, which establishes the claim. ■

References

- [1] Abadie, A., and G. W. Imbens (2006): “Large Sample Properties of Matching Estimators for Average Treatment Effects,” *Econometrica*, 74, 235 - 267.
- [2] Abadie, A., and G. W. Imbens (2011): “Bias-Corrected Matching Estimators for Average Treatment Effects,” *Journal of Business & Economic Statistics*, 29, 1 - 11.
- [3] Aitchison, J., and C. G. G. Aitken (1976): “Multivariate Binary Discrimination by the Kernel Method,” *Biometrika*, 63, 413 - 420.
- [4] Angrist, J., and A. Krueger (1992): “The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples,” *Journal of the American Statistical Association*, 87, 328 - 336.
- [5] Angrist, J., and A. Krueger (1995): “Split-Sample Instrumental Variables Estimates of the Return to Schooling,” *Journal of Business & Economic Statistics*, 13, 225 - 235.
- [6] Arellano, M., and C. Meghir (1992): “Female Labour Supply and on the Job Search: An Empirical Model Estimated Using Complementary Data Sets,” *Review of Economic Studies*, 59, 537 - 559.
- [7] Bekker, P. A. (1994): “Alternative Approximations to the Distributions of Instrumental Variable Estimators,” *Econometrica*, 62, 657 - 681.
- [8] Black, B., M. Trainor, and J. E. Spencer (1999): “Wage Protection Systems, Segregation and Gender Pay Inequalities: West Germany, the Netherlands and Great Britain,” *Cambridge Journal of Economics*, 23, 449 - 464.
- [9] Blackburn, M., and D. Neumark (1992): “Unobserved Ability, Efficiency Wages, and Interindustry Wage Differentials,” *Quarterly Journal of Economics*, 107, 1421 - 1436.
- [10] Bound, J., D. A. Jaeger, and R. M. Baker (1995): “Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak,” *Journal of the American Statistical Association*, 90, 443 - 450.
- [11] Callan, T., and C. Harmon (1999): “The Economic Returns to Schooling in Ireland,” *Labour Economics*, 6, 543 - 550.
- [12] Card, D. (1995): “Using Geographic Variation in College Proximity to Estimate the Return to Schooling,” in L. N. Christophides, E. K. Grant, and R. Swidinsky (eds.), *Aspects of Labour Market Behavior: Essays in Honour of John Vanderkamp*. Toronto: University of Toronto Press, 201 - 222.
- [13] Card, D. (2001): “Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems,” *Econometrica*, 69, 1127 - 1160.
- [14] Chao, J. C., and N. R. Swanson (2005): “Consistent Estimation with a Large Number of Weak Instruments,” *Econometrica*, 73, 1673 - 1692.
- [15] Chen, S. X. (1999): “Beta Kernel Estimators for Density Functions,” *Computational Statistics and Data Analysis*, 31, 131 - 145.
- [16] Choi, J., J. Gu, and S. Shen (2018): “Weak-Instrument Robust Inference for Two-Sample Instrumental Variable Regression,” *Journal of Applied Econometrics*, 33, 109 - 125.

- [17] Dieterle, S. G., and A. Snell (2016): “A Simple Diagnostic to Investigate Instrument Validity and Heterogeneous Effects When Using a Single Instrument,” *Labour Economics*, 42, 76- 86.
- [18] Eicker, F. (1963): “Asymptotic Normality and Consistency of the Least Squares Estimators for Families of Linear Regressions,” *Annals of Mathematical Statistics*, 34, 447- 456.
- [19] Fan, J., and Q. Yao (1998): “Efficient Estimation of Conditional Variance Functions in Stochastic Regression,” *Biometrika*, 85, 645- 660.
- [20] Fan, Y., and Q. Li (1996): “Consistent Model Specification Tests: Omitted Variables and Semiparametric Functional Forms,” *Econometrica*, 64, 865- 890.
- [21] Fang, H., M. P. Keane, and D. Silverman (2008): “Sources of Advantageous Selection: Evidence from the Medigap Insurance Market,” *Journal of Political Economy*, 116, 303- 350.
- [22] Flavin, M., and S. Nakagawa (2008): “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence,” *American Economic Review*, 98, 474- 495.
- [23] Flores-Lagunes, A. (2007): “Finite Sample Evidence of IV Estimators under Weak Instruments,” *Journal of Applied Econometrics*, 22, 677- 694.
- [24] Griliches, Z. (1977): “Estimating the Return to Schooling: Some Econometrics Problems,” *Econometrica*, 45, 1- 22.
- [25] Griliches, Z. (1979): “Sibling Models and Data in Economics: Beginnings of a Survey” *Journal of Political Economy*, LXXXVII, Part 2, S37-64.
- [26] Guerre, E., I. Perrigne, and Q. Vuong (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68, 525- 574.
- [27] Hansen, B. E. (2005): “Exact Mean Integrated Squared Error of Higher Order Kernel Estimators,” *Econometric Theory*, 21, 1031- 1057.
- [28] Hansen, B. E. (2008): “Uniform Convergence Rates for Kernel Estimation with Dependent Data,” *Econometric Theory*, 24, 726- 748.
- [29] Heckman, J. J., H. Ichimura, and P. Todd (1998): “Matching as an Econometric Evaluation Estimator,” *Review of Economic Studies*, 65, 261- 294.
- [30] Hirukawa, M., and A. Prokhorov (2018): “Consistent Estimation of Linear Regression Models Using Matched Data,” *Journal of Econometrics*, 203, 344- 358.
- [31] Hirukawa, M., Murtazashvili I., and A. Prokhorov (2020): “Uniform Convergence Rates for Nonparametric Estimators Smoothed by the Beta Kernel,” *Technical Report*.
- [32] Hoogerheide, L., J. Brock, and R. Thurik (2012): “Family Background Variables as Instruments for Education in Income Regressions: A Bayesian Analysis,” *Economics of Education Review*, 31, 515- 523.
- [33] Horowitz, J. L., and V. G. Spokoiny (2001): “An Adaptive, Rate-Optimal Test of a Parametric Mean-Regression Model Against a Nonparametric Alternative,” *Econometrica*, 69, 599- 631.

- [34] Hsu, Y.-C., J.-C. Liao, and E. S. Lin (2018): “Two-step Series Estimation and Specification Testing of (Partially) Linear Models with Nonparametrically Generated Regressors,” Working Paper, Academia Sinica.
- [35] Inoue, A., and G. Solon (2010): “Two-Sample Instrumental Variables Estimators,” *Review of Economics and Statistics*, 92, 557-561.
- [36] Klevmarcken, A. (1982): “Missing Variables and Two-Stage Least Squares Estimation from More than One Data Set,” IUI Working Paper, No. 62, The Research Institute of Industrial Economics (IUI), Stockholm.
- [37] Lavergne, P. (2001): “An Equality Test across Nonparametric Regressions,” *Journal of Econometrics*, 103, 307-344.
- [38] Li, Q., and D. Ouyang (2005): “Uniform Convergence Rate of Kernel Estimation with Mixed Categorical and Continuous Data,” *Economics Letters*, 86, 291-296.
- [39] Marron, J. S. (1994): “Visual Understanding of Higher-Order Kernels,” *Journal of Computational and Graphical Statistics*, 3, 447-458.
- [40] Marron, J. S., and M. P. Wand (1992): “Exact Mean Integrated Squared Error,” *Annals of Statistics*, 20, 712-736.
- [41] Mincer, J. A. (1974): *Schooling, Experience and Earnings*. New York: National Bureau of Economic Research.
- [42] Murtazashvili, I., D. Liu, and A. Prokhorov (2005): “Two-sample nonparametric estimation of intergenerational income mobility in the United States and Sweden,” *Canadian Journal of Economics*, 48, 1733-1761.
- [43] Newey, W. K. (1994): “Kernel Estimation of Partial Means and a General Variance Estimator,” *Econometric Theory*, 10, 233-253.
- [44] Newey, W. K. (2009): “Two-Step Series Estimation of Sample Selection Models,” *Econometrics Journal*, 12, S217-S229.
- [45] Pacini, D. (2019): “Two-Sample Least Squares Projection,” *Econometric Reviews*, 38, 95-123.
- [46] Pagan, A. (1984): “Econometric Issues in the Analysis of Regressions with Generated Regressors,” *International Economic Review*, 25, 221-247.
- [47] Parker, S. C., and C. M. Van Praag (2006): “Schooling, Capital Constraints, and Entrepreneurial Performance: The Endogenous Triangle,” *Journal of Business & Economic Statistics*, 24, 416-431.
- [48] Racine, J. S. (1997): “Consistent Significance Testing for Nonparametric Regression,” *Journal of Business & Economic Statistics*, 15, 369-378.
- [49] Racine, J. S., J. Hart, and Q. Li (2006): “Testing the Significance of Categorical Predictor Variables in Nonparametric Regression Models,” *Econometric Reviews*, 25, 523-544.
- [50] Racine, J. S., and Q. Li (2004): “Nonparametric Estimation of Regression Functions with Both Categorical and Continuous Data,” *Journal of Econometrics*, 119, 99-130.
- [51] Rilstone, P. (1996): “Nonparametric Estimation of Models with Generated Regressors,” *International Economic Review*, 37, 299-313.

- [52] Stengos, T., and B. Yan (2001): “Double Kernel Nonparametric Estimation in Semiparametric Econometric Models,” *Journal of Nonparametric Statistics*, 13, 883-906.
- [53] Stone, C. J. (1982): “Optimal Global Rates of Convergence for Nonparametric Regression,” *Annals of Statistics*, 10, 1040-1053.
- [54] Su, L., I. Murtazashvili, and A. Ullah (2013): “Local Linear GMM Estimation of Functional Coefficient IV Models with an Application to Estimating the Rate of Return to Schooling,” *Journal of Business & Economic Statistics*, 31, 184-207.
- [55] Veroff, J., L. McClelland, and K. Marquis (1971): “Measuring Intelligence and Achievement Motivation in Surveys,” *Technical Series Paper #71-01, Survey Research Center-Institute for Social Research, University of Michigan, Ann Arbor*.
- [56] White, H. (1980): “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity,” *Econometrica*, 48, 817-838.
- [57] Wooldridge, J. M. (2013): *Introductory Econometrics: A Modern Approach*, 5th Edition. Mason, OH: South-Western Cengage Learning.
- [58] Yatchew, A. (1997): “An Elementary Estimator of the Partial Linear Model,” *Economics Letters*, 57, 135-143.
- [59] Zabalza, A., and J.L. Arrufat (1985): “The Extent of Sex Discrimination in Great Britain,” in A. Zabalza and Z. Tzannatos (eds.), *Women and Equal Pay: The Effects of Legislation on Female Employment and Wages in Britain*, Cambridge, U.K.: Cambridge University Press, 70-96.

Table 1: Monte Carlo Results

Estimator	$(n, m) = (1000, 500)$					$(n, m) = (2000, 1000)$				
	Mean	SD	RMSE	SE	CR	Mean	SD	RMSE	SE	CR
Model A: $h(x) = x$										
OLS*	0.9998	0.0336	0.0336	0.0334	95%	0.9998	0.0236	0.0236	0.0235	95%
OLS-S	2.0004	0.0516	1.0017	—	—	1.9997	0.0383	1.0004	—	—
IV-S: $\rho = 0.4$	1.0004	0.1460	0.1460	0.1431	95%	1.0042	0.1019	0.1020	0.1008	94%
$\rho = 0.1$	0.9028	0.8699	0.8754	0.8039	95%	0.9376	0.5032	0.5070	0.4897	96%
MSOLS	1.3686	0.0639	0.3741	—	—	1.3684	0.0441	0.3711	—	—
MSII	0.9671	0.1105	0.1153	—	—	0.9868	0.0714	0.0726	—	—
MSII-FM	0.9518	0.1125	0.1224	0.0972	91%	0.9795	0.0719	0.0748	0.0656	93%
PARA	1.0004	0.0727	0.0727	0.0725	95%	1.0035	0.0533	0.0534	0.0514	94%
PILS-E	1.0847	0.0773	0.1147	0.0881	85%	1.0569	0.0541	0.0785	0.0650	89%
PILS-B	1.0080	0.0769	0.0773	0.0815	96%	1.0108	0.0549	0.0559	0.0602	96%
Model B: $h(x) = \sin(\pi x/4)$										
OLS*	0.9996	0.0295	0.0295	0.0297	95%	1.0000	0.0212	0.0212	0.0210	95%
OLS-S	1.6084	0.0517	0.6106	—	—	1.6075	0.0383	0.6087	—	—
IV-S: $\rho = 0.4$	1.0023	0.1391	0.1391	0.1367	96%	1.0045	0.0973	0.0974	0.0965	95%
$\rho = 0.1$	0.9459	0.7678	0.7697	0.7171	96%	0.9591	0.4442	0.4461	0.4417	97%
MSOLS	1.2205	0.0573	0.2278	—	—	1.2227	0.0399	0.2263	—	—
MSII	0.9800	0.0860	0.0883	—	—	0.9937	0.0565	0.0568	—	—
MSII-FM	0.9733	0.0869	0.0909	0.0678	86%	0.9906	0.0567	0.0575	0.0465	89%
PARA	0.9994	0.0620	0.0620	0.0621	95%	1.0029	0.0455	0.0456	0.0440	93%
PILS-E	1.0468	0.0674	0.0820	0.0717	91%	1.0330	0.0465	0.0571	0.0529	93%
PILS-B	0.9961	0.0668	0.0669	0.0664	96%	1.0019	0.0470	0.0471	0.0490	96%
Model C: $h(x) = \{(5/3)x + 3/2\} \mathbf{1}\{x < 0\} + \{-(4/3)x + 3/2\} \mathbf{1}\{x \geq 0\}$										
OLS*	0.9996	0.0266	0.0266	0.0276	96%	1.0001	0.0199	0.0199	0.0195	94%
OLS-S	1.1658	0.0520	0.1738	—	—	1.1667	0.0384	0.1711	—	—
IV-S: $\rho = 0.4$	1.0002	0.1314	0.1314	0.1334	95%	1.0041	0.0912	0.0913	0.0943	96%
$\rho = 0.1$	0.9855	0.6594	0.6595	0.6506	98%	1.0076	0.3991	0.3992	0.4120	97%
MSOLS	1.0601	0.0527	0.0799	—	—	1.0630	0.0381	0.0736	—	—
MSII	0.9913	0.0695	0.0700	—	—	0.9996	0.0472	0.0472	—	—
MSII-FM	0.9896	0.0699	0.0706	0.0502	84%	0.9988	0.0473	0.0473	0.0350	85%
PARA	0.9561	0.0825	0.0935	—	—	0.9586	0.0594	0.0724	—	—
PILS-E	1.0111	0.0607	0.0617	0.0710	97%	1.0104	0.0419	0.0432	0.0532	98%
PILS-B	0.9959	0.0604	0.0605	0.0664	97%	1.0008	0.0420	0.0420	0.0495	98%
Model D: $h(x) = x + (5/\tau)\phi(x/\tau) - (5/2)\{\Phi(2/\tau) - 1/2\}$, $\tau = 3/4$										
OLS*	1.0004	0.0325	0.0325	0.0332	95%	0.9997	0.0233	0.0233	0.0235	95%
OLS-S	1.9992	0.0513	1.0005	—	—	2.0000	0.0379	1.0008	—	—
IV-S: $\rho = 0.4$	0.9966	0.1419	0.1419	0.1432	96%	1.0038	0.1013	0.1013	0.1008	94%
$\rho = 0.1$	0.8913	0.8477	0.8547	0.8034	95%	0.9587	0.4982	0.4999	0.4896	96%
MSOLS	1.3634	0.0624	0.3687	—	—	1.3616	0.0445	0.3644	—	—
MSII	0.9556	0.1089	0.1176	—	—	0.9803	0.0682	0.0710	—	—
MSII-FM	0.9397	0.1113	0.1266	0.0985	90%	0.9730	0.0689	0.0740	0.0655	94%
PARA	0.7369	0.1320	0.2944	—	—	0.7371	0.0948	0.2795	—	—
PILS-E	1.0774	0.0752	0.1079	0.0884	89%	1.0510	0.0519	0.0728	0.0652	91%
PILS-B	0.9884	0.0761	0.0770	0.0830	96%	0.9959	0.0529	0.0531	0.0609	98%

Notes: *Mean* = simulation average of the parameter estimate; *SD* = simulation standard deviation of the parameter estimate; *RMSE* = root mean-squared error of the parameter estimate; *SE* = simulation average of the standard error; and *CR* = coverage rate for the nominal 95% confidence interval.

Table 2: **Sample Characteristics**

Variable	Mean	Std. Dev.	Min.	Median	Max.
Panel A: The PSID Sample					
<i>Earnings</i>	8763.24	5968.39	30	7900	70000
$\log(\textit{Earnings})$	8.84	0.80	3.40	8.97	11.16
<i>Education</i>	12.15	3.09	5	12	17
<i>Experience</i>	19.57	13.17	0	18	68
<i>Married</i>	0.89	0.31	0	1	1
<i>Black</i>	0.26	0.44	0	0	1
<i>South</i>	0.41	0.49	0	0	1
<i>Urban</i>	0.28	0.45	0	0	1
<i>Age</i>	38.57	13.48	17	37	86
<i>South while growing up</i>	0.46	0.50	0	0	1
<i>IQ score</i>	9.49	2.28	0	10	13
<i>Father's education</i>	9.26	3.06	0	8	17
Panel B: The NLS sample					
<i>Education</i>	12.98	3.20	1	13	18
<i>Married</i>	0.63	0.48	0	1	1
<i>Black</i>	0.34	0.48	0	0	1
<i>South</i>	0.53	0.50	0	1	1
<i>Urban</i>	0.58	0.49	0	1	1
<i>Age</i>	28.47	3.14	24	28	34
<i>South while growing up</i>	0.55	0.50	0	1	1
<i>KWW score</i>	32.48	8.42	4	33	56

Notes: ($n =$)2430 individuals in the PSID sample. ($m =$)1102 individuals in the NLS sample. Demeaned *IQ score* and *KWW score* are used.

Table 3: Estimation Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS*	OLS-S	IV-S	MSOLS	MSII-FM	PARA	PILS
<i>Education</i>	0.0635 (0.0056)	0.0718 (0.0053)	0.0953 (0.0159)	0.0727 (0.0059)	0.0685 (0.0080)	-0.2850 (0.2281)	0.0568 (0.0087)
<i>Experience</i>	0.0809 (0.0043)	0.0818 (0.0043)	0.0830 (0.0045)	0.0826 (0.0049)	0.0766 (0.0065)	0.1170 (0.0095)	0.0818 (0.0043)
<i>Experience</i> ²	-0.0017 (0.0001)	-0.0017 (0.0001)	-0.0017 (0.0001)	-0.0017 (0.0001)	-0.0016 (0.0001)	-0.0017 (0.0001)	-0.0017 (0.0001)
<i>Ability</i>	0.0313 (0.0078)	- (-)	- (-)	-0.0012 (0.0027)	0.0029 (0.0081)	0.2444 (0.1514)	0.0094 (0.0045)
<i>Married</i>	0.3717 (0.0539)	0.3793 (0.0536)	0.3844 (0.0536)	0.3799 (0.0535)	0.3777 (0.0535)	0.2226 (0.1319)	0.3958 (0.0543)
<i>Black</i>	-0.1302 (0.0323)	-0.1741 (0.0316)	-0.1249 (0.0426)	-0.1849 (0.0393)	-0.1504 (0.0806)	-0.1819 (0.1146)	-0.1776 (0.0316)
<i>South</i>	-0.0921 (0.0288)	-0.0983 (0.0287)	-0.0814 (0.0314)	-0.0979 (0.0286)	-0.0989 (0.0289)	-0.1372 (0.1135)	-0.1042 (0.0287)
<i>Urban</i>	0.1363 (0.0282)	0.1499 (0.0284)	0.1278 (0.0332)	0.1538 (0.0297)	0.1404 (0.0390)	0.0138 (0.1273)	0.1557 (0.0282)
Data Combination?	No	No	No	Yes	Yes	Yes	Yes
Sample Size: <i>n</i>	2430	2430	2430	2430	2430	2430	2430
<i>m</i>	-	-	-	1102	1102	1102	1102

Notes: The dependent variable is the log of total annual labor earnings. *Education*, *married*, *black*, *south*, and *urban* are used as the included common variables, and *age* and *south while growing up* are used as the excluded common variables. The (demeaned) *IQ score* and *KWW score* variables are used as ability measures in the PSID and NLS samples, respectively.